

What is claimed is:

1. A method for arbitrating for a switch fabric having a plurality of ports, each port from the plurality of ports having its own plurality of links, comprising:
 - determining, on a per port basis, a subset of links from the plurality of links associated with that port, each link from the determined subset of links for that port
 - 5 being associated with a candidate packet, each link from the plurality of links for that port being associated with a weight value; and
 - selecting, on a per port basis, a link from the determined subset of links for that port based on the weight value for determined subset of links for that port.
2. The method of claim 1, wherein the weight value for the selected link for
- 10 that port is greater than the weight value associated with each remaining link from the determined subset of links for that port.
3. The method of claim 1, wherein the weight value associated with each link from the plurality of links associated with that port corresponds to a priority for that port and a second port associated with that link.
- 15 4. The method of claim 1, wherein the candidate packet associated with a link from the plurality of links is buffered at an input port associated with that link.
5. The method of claim 1, wherein the determining and the selecting are performed on a per output port basis.
6. The method of claim 5, further comprising:
 - 20 receiving, on a per output port basis, a plurality of requests associated with that port, each request from the plurality of requests associated with that port being associated with a link from the plurality of links associated with that port and being associated with a candidate packet; and
 - sending, on a per output port basis, a grant associated with the selected link for
 - 25 that output port.

7. The method of claim 6, further comprising:
determining, on a per input port basis, a subset of links from a plurality of links associated with that input port based on any grants received at that input port; and
selecting, on a per input port basis, at most one link from the determined subset
5 of links associated with that input port and having a weight value greater than a weight value associated with each remaining link from the determined subset of links for that input port.
8. The method of claim 1, wherein the determining and the selecting are performed on a per input port basis.
- 10 9. The method of claim 8, further comprising:
receiving, on a per input port basis, a plurality of requests associated with that port, each request from the plurality of requests associated with that port being associated with a link from the plurality of links associated with that port and being associated with a candidate packet; and
15 sending, on a per input port basis, a grant associated with the selected link for that input port.
10. The method of claim 9, further comprising:
determining, on a per output port basis, a subset of links from a plurality of links associated with that output port based on any grants received at that output port;
20 and
selecting, on a per output port basis, at most one link from the determined subset of links associated with that output port and having a weight value greater than a weight value associated with each remaining link from the determined subset of links for that output port.
- 25 11. A method for arbitrating for a switch fabric, comprising:
determining, for a first port, a link subset from a plurality of links associated with the first port, each link from the link subset being associated with its own candidate packet and being associated with its own weight value;

selecting a link from the link subset for the first port based on the weight value associated with each link from the link subset for the first port;

determining, for a second port, a link subset from a plurality of links associated with the second port, each link from the link subset associated with the second port
5 being associated with its own candidate packet and being associated with its own weight value, the determining for the second port being performed in parallel with the determining for the first port; and

selecting a link from the link subset for the second port based on the weight value associated with each link from the link subset of associated with the second port,
10 the selecting for the second port being performed in parallel with the selecting for the first port.

12. The method of claim 11, wherein the weight value associated with each link from the plurality of links associated with the first port corresponds to a priority between the first port and a remaining port associated with that link.

15 13. The method of claim 11, further comprising:
sending, from the first port, a grant associated with link selected for the first port;

sending, from the second port, a grant associated with the link selected for the second port;

20 determining, for a third port, a subset of links from a plurality of links associated with the third port based on any grants received at the third port; and

selecting, for the third port, at most one link from the determined subset of links associated with the third port based on the weight value associated with each link from the determined subset of links for the third port.

25 14. The method of claim 11, wherein the first port is an output port and the second port is an output port.

15. The method of claim 14, further comprising:
receiving, at the first port, a plurality of requests associated with the first port, each request from the plurality of requests associated with the first port being associated with a candidate packet;
5 sending, from the first port, a grant associated with the selected link for first port;
receiving, at the second port, a plurality of requests associated with the second port, each request from the plurality of requests associated with the second port being associated with a candidate packet; and
10 sending, from the second port, a grant associated with the selected link for second port.
16. The method of claim 15, further comprising:
determining, for a first input port, a subset of links from a plurality of links associated with the first input port based on any grants received at the first input port;
15 and
selecting, for a first input port, at most one link from the determined subset of links associated with the first input port and having a weight value greater than a weight value associated with each remaining link from the determined subset of links for the first input port.
- 20 17. The method of claim 11, wherein the first port is an input port and the second port is an input port.
18. The method of claim 17, further comprising:
receiving, at the first port, a plurality of requests associated with the first port, each request from the plurality of requests associated with the first port being
25 associated with a candidate packet;
sending, from the first port, a grant associated with the selected link for first port;
receiving, at the second port, a plurality of requests associated with the second port, each request from the plurality of requests associated with the second port being
30 associated with a candidate packet; and

sending, from the second port, a grant associated with the selected link for second port.

19. The method of claim 18, further comprising:

5 determining, for a first output port, a subset of links from a plurality of links associated with the first output port based on any grants received at the first output port; and

10 selecting, for a first output port, at most one link from the determined subset of links associated with the first output port and having a weight value greater than a weight value associated with each remaining link from the determined subset of links for the first output port.

20. An apparatus, comprising:

a plurality of input ports;

a plurality of output ports each being coupled to each input port from the plurality of input ports by a link; and

15 a plurality of arbiters each being uniquely associated with its own output port from the plurality of output ports, the plurality of arbiters being configured to operate in parallel with each other, a first arbiter from the plurality of arbiters being associated with a first output port from the plurality of output ports, the first arbiter having a selection unit configured to transmit an arbitration signal based on a weight value
20 associated with each link associated with the first output port and having a candidate packet.

21. The apparatus of claim 20, further comprising:

an arbiter associated with a first input port from the plurality of input ports and having a selection unit configured to receive at least one arbitration signal from a
25 respective output port from the plurality of output ports, the selection unit of the arbiter associated with the first input port being configured to transmit an arbitration signal based on a weight value associated with each link associated with the received at least one arbitration signal.

22. The apparatus of claim 21, further comprising:

a second plurality of arbiters each being uniquely associated with its own input port from the plurality of input ports, the second plurality of arbiters being configured to operate in parallel with each other.

5 23. The apparatus of claim 22, wherein:

each arbiter from the second plurality of arbiters is configured to receive the arbitration signal from any arbiter from the first plurality of arbiters, an arbiter associated with a first input port from the plurality of input ports being configured to transmit an arbitration signal based on: (1) a weight value associated with each link
10 associated with the first input port and (2) any arbitration signals received from the plurality of arbiters.

24. An apparatus, comprising:

a plurality of input ports;

15 a plurality of output ports each being coupled to each input port from the plurality of input ports by a link; and

a plurality of arbiters each being uniquely associated with its own input port from the plurality of input ports, the plurality of arbiters being configured to operate in parallel with each other, a first arbiter from the plurality of arbiters being associated with a first input port from the plurality of input ports, the first arbiter having a
20 selection unit configured to transmit an arbitration signal based on a weight value associated with each link associated with the first input port and having a candidate packet.

25. The apparatus of claim 24, further comprising:

25 an arbiter associated with a first output port from the plurality of output ports and having a selection unit configured to receive at least one arbitration signal from a respective input port from the plurality of input ports, the selection unit of the arbiter associated with the first output port being configured to transmit an arbitration signal based on a weight value associated with each link associated with the received at least one arbitration signal.

26. The apparatus of claim 24, further comprising:

a second plurality of arbiters each being uniquely associated with its own output port from the plurality of output ports, the second plurality of arbiters being configured to operate in parallel with each other.

5 27. The apparatus of claim 26, wherein:

each arbiter from the second plurality of arbiters is configured to receive the arbitration signal from any arbiter from the first plurality of arbiters, an arbiter associated with a first output port from the plurality of output ports being configured to transmit an arbitration signal based on: (1) a weight value associated with each link
10 associated with the first output port and (2) any arbitration signals received from the plurality of arbiters.

28. An apparatus, comprising:

a first selection unit associated with its own plurality of links, the first selection unit being configured to receive an activate-link signal and being configured to transmit
15 an arbitration signal based on a weight value associated with each link from the plurality of links for the first selection unit and based on the activation-link signal received at the first selection unit;

a second selection unit associated with its own plurality of links, the second selection unit being configured to receive an activate-link signal and being configured
20 to transmit an arbitration signal based on a weight value associated with each link from the plurality of links for the second selection unit and based on the activation-link signal received at the second selection unit; and

a third selection unit associated with its own plurality of links, the third selection unit being configured to receive any arbitration signals associated with the
25 plurality of links for the third selection unit, the third selection unit being configured to transmit an arbitration signal based on a weight value associated with each link from the plurality of links for the third selection unit and based on any arbitration signals received at the third selection unit.

29. The apparatus of claim 28, wherein:

the first selection unit is associated with a first output port, each link from the plurality of links for the first selection unit couples the first output port to an input port from a plurality of input ports, the arbitration signal for the first selection unit indicates a grant associated with the first output port and with an input port from the plurality of input ports; and

the second selection unit is associated with a second output port, each link from the plurality of links for the second selection unit couples the second output port to an input port from a plurality of input ports, the arbitration signal for the second selection unit indicates a grant associated with the second output port and with an input port from the plurality of input ports.

30. The apparatus of claim 29, wherein:

the third selection unit is associated with a first input port from the plurality of input ports, the arbitration signal to be transmitted from the third selection unit indicates an accept associated with the first input port and any output ports associated with any respective arbitration signal received at the third selection unit.

31. The apparatus of claim 28, wherein:

the first selection unit is associated with a first input port, each link from the plurality of links for the first selection unit couples the first input port to an output port from a plurality of output ports, the arbitration signal for the first selection unit indicates a grant associated with the first input port and with an output port from the plurality of output ports; and

the second selection unit is associated with a second input port, each link from the plurality of links for the second selection unit couples the second port to an output port from a plurality of output ports, the arbitration signal for the second selection unit indicates a grant associated with the second input port and with an output port from the plurality of output ports.

32. The apparatus of claim 31, further comprising:

the third selection unit is associated with a first output port from the plurality of output ports, the arbitration signal to be transmitted from the third selection unit

indicates an accept associated with the first output port and any input ports associated with any respective arbitration signal received at the third selection unit.

33. The apparatus of claim 28, wherein:
 the first selection unit is configured to operate in parallel with the second
5 selection unit; and
 the third selection unit is configured to operate in parallel with a fourth selection
 unit.
34. A method for arbitrating for a switch, comprising:
 determining, in a first time slot, a plurality of grants associated with a plurality
10 of output arbiters; and
 determining, in the first time slot, a plurality of grants associated with a
 plurality of input arbiters.
35. The method of claim 34, further comprising:
 determining, in a second time slot, a plurality of accepts associated with the
15 plurality of input arbiters based on the plurality of grants determined for the plurality of
 output arbiters in the first time slot; and
 determining, in the second time slot, a plurality of accepts associated with the
 plurality of output arbiters based on the plurality of grants determined for the plurality
 of input arbiters in the first time slot.
- 20 36. The method of claim 35, further comprising:
 generating a plurality of arbitration selections based on the plurality of accepts
 associated with the plurality of input arbiters and based on the plurality of accepts
 associated with the plurality of output arbiters.
37. The method of claim 35, further comprising:
25 selecting one from the group of (1) the plurality of accepts associated with the
 plurality of input arbiters, and (2) the plurality of accepts associated with the plurality
 of output arbiters, based on an optimal throughput associated with that selected group.
38. The method of claim 35, further comprising:

selecting a plurality of arbitration selections based on the plurality of accepts associated with the plurality of input arbiters and based on the plurality of accepts associated with the plurality of output arbiters, the plurality of arbitration selections having a first portion and a second portion, the first portion of the plurality of arbitration selections being associated with a portion from the plurality of accepts associated with the plurality of input arbiters, the second portion of the plurality of arbitration selections being associated with a portion from the plurality of accepts associated with the plurality of output arbiters.

39. A method for arbitrating for a switch, comprising:
 10 determining a grant for a first output arbiter from a plurality of output arbiters;
 determining a grant for a first input arbiter from a plurality of input arbiters, the determining the grant for the first input arbiter being performed in parallel with the determining the grant for the first output arbiter;
 determining an accept for a second input arbiter from the plurality of input
 15 arbiters based on the determined grant for the first output arbiter; and
 determining an accept for a second output arbiter from the plurality of output arbiters based on the determined grant for the first input arbiter, the determining the accept for second output arbiter being performed in parallel with the determining the accept for the second input arbiter.

20 40. The method of claim 39, further comprising:
 selecting one from the group of: (1) the accept for the second input arbiter and
 (2) the accept for the second output arbiter.

41. The method of claim 39, further comprising:
 selecting, for a first time slot, one from the group of: (1) the accept for the
 25 second input arbiter and (2) the accept for the second output arbiter, the determining the accept for the second input arbiter and the accept for the second output arbiter being performed for the first time slot; and
 selecting, for a second time slot, one from the group of: (1) an accept for the second input arbiter for the second time slot and (2) an accept for the second output
 30 arbiter for the second time slot.

42. A method for arbitrating for a switch fabric, comprising:

determining, for a first output port from a plurality of output ports, a link from a plurality of links associated with the first output port based on a weight value associated with each link from a subset of links from the plurality of links, each link from the subset of links being associated with its own candidate packet; and

determining, for a first input port from a plurality of input ports, a link from a plurality of links associated with the first input port based on a weight value associated with each link from a subset of links from the plurality of links, each link from the subset of links being associated with its own candidate packet, the determining for the first input port being performed in parallel with the determining for the first output port.

43. The method of claim 42, further comprising:

selecting, for a second output port from the plurality of output ports, a link from a plurality of links associated with the second output port including the link determined for the first input port, the selecting for the second output being based on a weight value associated with each link from the plurality of links associated with the second output port; and

selecting, for a second input port from the plurality of input ports, a link from a plurality of links associated with the second input port including the link determined for the first output port, the selecting for the second input being based on a weight value associated with each link from the plurality of links associated with the second input port, the selecting for the second input port being performed in parallel with the selecting for the second output port.

44. The method of claim 42, wherein:

the link selected for the second output port is selected from a subset of links from plurality of links for the second output port, each link from the subset of links associated with the second output port is associated with a determination of a corresponding input port; and

the link selected for the second input port is selected from a subset of links from plurality of links for the second input port, each link from the subset of links associated with the second input port is associated with a determination of a corresponding output port.

45. An apparatus, comprising:
a plurality of first-stage arbiters; and
a plurality of second-stage arbiters each associated with the plurality of first-stage arbiters,
- 5 each first-stage arbiter from the plurality of first-stage arbiters being configured to send its own grant signal to a second-stage arbiter from the plurality of second-stage arbiters during a first time slot, each second-stage arbiter from the plurality of second-stage arbiters being configured to send its own grant signal to a first-stage arbiter from the plurality of first-stage arbiters during the first time slot.
- 10 46. The apparatus of claim 46, wherein:
a first first-stage arbiter from the plurality of first-stage arbiters being configured to receive at least one grant from a second-stage arbiter from the plurality of second-stage arbiters during a second time slot, the first first-stage arbiter being configured to send an accept signal based on any grant received from a second-stage
- 15 arbiter from the plurality of second-stage arbiters; and
a first second-stage arbiter from the plurality of second-stage arbiters being configured to receive at least one grant from a first-stage arbiter from the plurality of first-stage arbiters during the second time slot, the first second-stage arbiter being configured to send an accept signal based on any grant received from a first-stage
- 20 arbiter from the plurality of first-stage arbiters.
47. The apparatus of claim 46, further comprising:
a matching combiner coupled to the plurality of first-stage arbiters and the plurality of second-stage arbiters, the matching combiner configured to receive a plurality of accept signals from the plurality first-stage arbiters and to receive a
- 25 plurality of accept signals from the plurality of second-stage arbiters, the matching combiner configured to send a plurality of arbitration selections based on the plurality of accept signals from the plurality first-stage arbiters and the plurality of accept signals from the plurality of second-stage arbiters.